



# Optical effects on HDR calibration via a multiple exposure noise-based workflow

Brian A. Karr<sup>1,2</sup> · Kurt Debattista<sup>1</sup> · Alan G. Chalmers<sup>1</sup>

Published online: 18 April 2020  
© The Author(s) 2020

## Abstract

High dynamic range (HDR) technology allows more of the lighting in a specific scene to be captured at a set point in time, and thus is capable of delivering an overall view of the scene that more closely correlates with our visual experience in the real world, compared to standard, or low dynamic range (LDR) technology. Although HDR capabilities of single exposure capture systems are improving, the traditional method for creating HDR images still includes combing a number of different exposures, captured with an LDR system, into a single HDR image. Several use cases requiring absolute calibration of the resulting HDR luminance map have been undertaken, but none of these have provided a detailed analysis of the optical effects of glare on the results. We develop a calibrated HDR radiance map, including methodical linearization of captured image data, while characterizing the limitations due to the effects of optical glare. A purposely designed controlled test scene is used to challenge the calibrated reconstruction efforts, including low luminance levels, spatial inclusion of lens vignette over the full imaged area, and optical glare. Results demonstrate that even with careful processing and recombination of the LDR data, radiometric accuracy is limited as a result of glare. The proposed approach performs better than calibration methods in commercially available HDR recombination software.

**Keywords** High dynamic range · Veiling glare · Lens flare, vignette · Multiple exposure recombination · Radiometric calibration · Photon transfer curve · Flat field

## 1 Introduction

There is a wide range of light and color in the real world. While the human eye is capable, through adaption, of seeing scenes from starlight ( $10^{-5}$  cd/m<sup>2</sup>) to bright sunshine ( $10^6$  cd/m<sup>2</sup>), traditional imaging technology is not. Such Low Dynamic Range (LDR) technology is only able to capture light in a scene where the difference between the darkest and brightest regions, the so called dynamic range, is approximately 256:1. This can result in over- or under-exposed areas of the captured image in which substantial information is missing. High dynamic range (HDR) imaging has been developed to capture important differences in luminance at the same moment, and in the same scene and lighting con-


ditions. It is important to note however that HDR imaging may not be able to acquire all possible changes of luminance in different scenes. Additionally, radiometric calibration may be impeded due to glare related scene and lighting dependent effects.

For visual computing cases where realism is a must, for example to provide robust visibility in autonomous vehicles, or when making physically accurate light measurement where absolute precision is required, carefully calibrated HDR is fundamentally important. Additional applications for calibrated HDR imagery include gas analysis [1], rocket plume visualization [2–4], lighting analysis [5,6], and image processing [7,8]. Absolute calibration and determination of true light levels can also assist in the shaping of data requirements for imaging equipment. As an example, Fig. 1 includes different exposures of a frame containing a dynamic scene from a test flight of an experimental NASA launch and landing test vehicle developed under the project name Morpheus [9]. While the Morpheus vehicle ultimately resulted in 13 successful free flights, during an initial test flight the vehicle crashed, resulting in the oxygen tank over-pressurization

✉ Brian A. Karr  
b.karr@warwick.ac.uk  
<http://www.RockledgeDesignGroup.com>

<sup>1</sup> Warwick Manufacturing Group, University of Warwick, Coventry, UK

<sup>2</sup> Rockledge Design Group, Inc, Rockledge, USA



**Fig. 1** Example exposures illustrating overall dynamic range of an oxygen tank over-pressurization event

and explosion shown in Fig. 1. The capture device, an ARRI Alexa, is a professional level digital camera with reported dynamic range capability of 14 stops [10]. Figure 1a illustrates the sample Alexa frame, post-processed for display on a typical LDR monitor. With the exposure optimized per the pre-explosion average luminance level of the scene, the oxygen cloud is saturated with mostly white pixels. The raw data however contains additional luminance detail in the oxygen cloud, and when re-processed, shifting the displayed exposure range towards the highlights, the oxygen cloud detail is then observed as shown in Fig. 1b. In the third representation of the frame, Fig. 1c, the exposure range is further shifted, such that detail in the brightest area of the oxygen cloud can be seen.

What is not immediately evident from the images is how much of the available dynamic range of the imager is utilized? If the scene were initially darker, would an imager with even greater dynamic range be required? These questions lead to key insights, including for budgetary considerations, such as how much dynamic range is truly required, at what sensitivity, and for what additional cost? We develop a calibrated HDR radiance map to attempt to answer these questions, including methodical linearization of captured image data, while characterizing the limitations due to the effects of optical glare.

In this visual computing application, focus is primarily on a linear workflow, including the removal of non-linear contributions of the imaging system to obtain linear data for further processing. On the contrary, HDR in commercial, broadcast, and cinema markets focuses heavily on log-based workflows where HDR imagery is tone-mapped for display for humans within the limits of the available hardware and software [11]. The workflows are not mutually exclusive. The methods presented in this article carefully categorize error in a multiple exposure HDR recombination workflow via methodical calibration; including determination of sensor noise attributes, system gain response, and optical effects, such that the linear dynamic range of the imaging system can be characterized to a greater extent than previous methods.

The two primary steps of the multiple exposure recombination technique are (1) linearization of the data of each exposure by applying the inverse Camera Response Func-

tion (CRF), and (2) the generation of the HDR luminance map from a weighted average of scaled pixel values of the multiple exposures. Weighting schemes include those that are proportional to the Digital Number (DN) [12], those that are proportional to the CRF [13], and noise-based models [14]. Recent reviews of weighting schemes can be found in Granados et al. [15] and Griffiths et al. [16]. Noise-based approaches consider the individual noise sources of the acquisition process, giving higher weights to values containing lower variances. One limiting factor to obtain good dynamic range is noise in the dark areas of the captured raw digital image. This can be optimized both in the selection of the weighting factor [15–17], as well as in the capture strategy [17,18]. Properly accounting for the noise sources results in a calibrated linear camera model that minimizes uncertainty while producing more accurate results; an important factor if the final HDR image is to be used as a physically accurate irradiance representation.

Optical effects that influence image capture include vignette, shading, interference fringing, lens flare, and veiling glare. Vignette, shading, and interference fringing are classified as fixed pattern noise (FPN) sources, which can be removed by the process of flat fielding [19]. Lens flare and veiling glare are scene dependent effects that can reduce contrast visibility due to light scattering, and set a physical limit on the luminance range that can be accurately measured [20]. Lens flare manifests as sharp image artifacts resulting from hard edges, internal to the lens, such as the lens aperture and leaves [21]. At further distances lens flare becomes blurred, and the effects of stray light from lens surfaces, camera walls, and the sensor surface itself combine as veiling glare. Examples of veiling glare include internal multiple reflections between lens surfaces, scatter from lens elements due to imperfections or optical cements, and reflections from lens edges, mounts, and shutter blades [22].

The main contribution of this work is the detailed evaluation and measurement of optical-based glare on a calibrated multiple exposure LDR to HDR pipeline, illustrating substantial improvements over non-calibrated workflows of commercially available software. Other contributions of this article are:

- An HDR workflow for processing raw linear camera data from a Bayer-based sensor, using floating point data types throughout to maintain precision.
- Absolute calibration of an HDR luminance map.
- Characterization of optical-based glare and the resulting impact on dynamic range.

## 2 Background and related work

In this section we introduce self-calibrating procedures that have been developed for the purpose of combining multiple exposures into HDR radiance maps [13,23–25], as well as the concept of weighted averages. Generally, the film or digital camera response curve is first recovered, and the HDR radiance map is then determined from a weighted average of scaled pixel values from multiple exposures. One of the earliest proposed weighted averages by Mann and Picard [12] fitted the exposure data to a power curve for use as the weighting function. Debevec and Malik [23] introduced a self-calibrating procedure utilizing a physical property of imaging systems called reciprocity. Reciprocity specifies that in the ideal case, the exposure  $X$  is a product only of the irradiance  $E$  and the exposure time  $\Delta T$ . Their hat-based function assigns higher weights to middle exposure values [23]. Additional methods for implementing weighting functions have since been developed by Robertson et al. [26] (similar to Gaussian curve), Mitsunaga et al. [13] (first order approx. to signal to noise ratio), and Ward et al. [27] and Akyuz et al. [8] (modification of Mitsunaga implementing broad-hat model). Noise-based models for weighting functions have been proposed by Tsin et al. [14] (statistical characterization), Kirk et al. [28], Granados et al. [15] and Hasinoff et al. [17]. Noise-based models consider a more in-depth analysis of the acquisition process, including both temporal and spatial noise sources. The noise-based approach is increasingly important if the final HDR image is to be used for physically accurate irradiance representation in addition to visualization. Use of the software package Photosphere [29] to determine the camera response function and fuse multiple images was carried out by Inanici et al. [24], including determination of the vignette effect and Point Spread Function (PSF). Popadic et al. [30] present a comparison of four image fusion algorithms where the images are treated from a global approach as opposed to a traditional pixel level method. Finally, Griffiths et al. [16] develop an unbiased test methodology for determining the aggregated error while using weighted averages as a maximum likelihood estimator.

A benefit of a noise-based model for determining weighting averages is the consideration of individual noise sources of the sensor acquisition process. Early weighting methods for linear sensors provide variance estimates based on a simplified camera noise model, transferring the measurement

uncertainty into the weighting function. Granados et al. [15] establishes an optimal weighting function implementing a more rigorous noise model, one that addresses both the temporal and spatial noise sources. This is the method we have used as the foundation of our new approach, building on the method by performing a radiometric calibration. Granados' method implements an iterative optimization for the maximum likelihood estimate of the irradiance and its uncertainty, such that the weighting function is not perturbed by measurement noise. The predicted uncertainty can also be used to optimally denoise the resulting irradiance map. The noise-based optimal weighting function  $w_{\text{opt}}$  for HDR reconstruction is given as:

$$w_{\text{opt}}(v_i) = \frac{1}{\sigma_{X_i}^2} = \frac{t_i^2 g^2 a_j^2}{g^2 t_i (a_j \mu_x + 2\mu_D) + 2\sigma_R^2}. \quad (1)$$

where  $v_i$  is the non-saturated mean radiance of the  $i$ th exposure,  $\sigma_{X_i}^2$  the uncertainty of irradiance  $X_i$ ,  $t_i$  the exposure time,  $g$  the overall camera gain factor,  $a_j$  the per-pixel gain factor for pixel  $j$ ,  $\mu_x$  the mean irradiance,  $\mu_D$  the mean dark current, and  $\sigma_R^2$  the spatial variance of the readout noise.  $X_i$  and  $\sigma_{X_i}^2$  are valid for  $v < v_{\text{sat}}$ , where  $v_{\text{sat}}$  is the saturation limit. As part of the optimal recovery of  $\mu_x$ , estimates for the temporal readout noise parameters including spatial mean  $\mu_R$  and spatial variance  $\sigma_R^2$ , the saturation limit  $v_{\text{sat}}$ , and the overall camera gain  $g$ , are generated by calibrating the camera sensor using the photon transfer curve (PTC) method. The per-pixel gain factors  $a_j$ , including spatial noise components photo-response non-uniformity (PRNU) and dark current non-uniformity (DCNU), are estimated via flat field measurement with background subtraction.

With the PTC derived values, the per-pixel gain factors, and the LDR multiple exposure images  $v_i$  and dark frames  $b_i$ , the optimal HDR recombination based on a rigorous noise model is performed using iterative estimates of the irradiance and its uncertainty. The goal is to reconstruct the mean radiance having the lowest variance from the set of multiple exposures to obtain the maximum likelihood estimate per Granados (Eq. 1) [15].

## 3 Calibrating a multiple exposure HDR workflow

In this section we introduce our overall method of generating a calibrated HDR image from multiple exposures (illustrated in Fig. 5), and describe the new steps taken to extend Granados' noise-based recombination method to achieve a calibrated HDR radiance map. The availability of linear data from the captured image file is device dependent, generally requiring characterization of in-camera processing and opti-

cal effects to produce a CRF. The inverse CRF can then be applied, linearizing the data. Determination of the CRF, and estimation of the physically accurate scene irradiance, is effectively the calculation of the photometric mean irradiance for each pixel. The nonlinear camera response function  $f$  is determined via radiometric calibration that includes the nonlinear contributions of the camera capture system such that  $f = g^{-1}$ , where  $g$  is the inverse function mapping scaled image irradiance to pixel values. Although previous efforts to produce a calibrated radiance map [13,23–25] have included calibrations for many of the system non-linearity's and temporal/spatial noise, inverse functions have not yet accounted for glare. Further, due to both the scene and lighting specific spatial nature of glare, the effect will constantly change. Our goal is to perform a first step of characterizing the glare effect for a static scene. The calibration steps, to be described below, include:

- Use of the PTC to determine the noise-based weighted average performance parameters: read noise, saturation, and camera gain.
- Determination of the per-pixel gain factors by flat field measurement and background subtraction, implementing an integrating sphere for the measurements.
- Radiometric scene calibration, and effect of glare in a static scene, by measurement and comparison to absolute luminance spot meter readings.

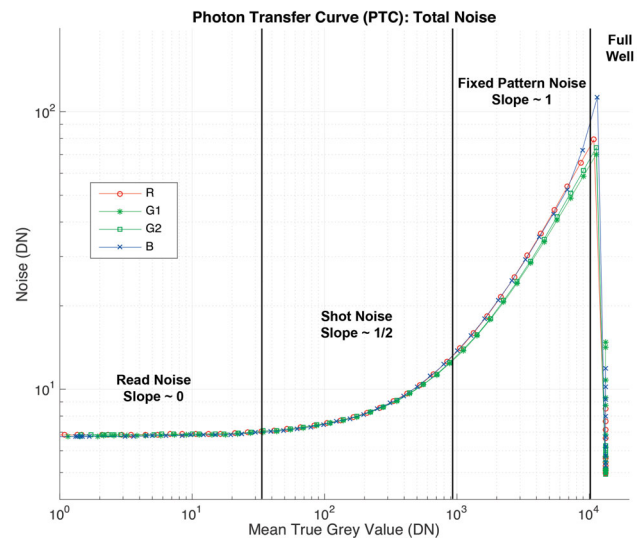
### 3.1 Performance parameters via PTC

To characterize error for the benefit of determining the ideal low noise pixels to be used in Granados's multiple exposure recombination process, we implement photon transfer method as a means of determining sensor performance parameters. PTCs are useful in the determination of camera system response to a uniform light source without optics. The performance parameters required for our noise model, including read noise mean and variance, charge capacity, and incident photon to ADC sensitivity (camera gain), can be determined from the PTC. The PTC is measured using an integrating sphere setup following the EMVA Standard 1288 [31], as shown in Fig. 2.

The PTC contains four identifiable regions dominated by a particular noise type, as described by Janesick [19]. An example for the Canon 5DM3 is shown in Fig. 3. The first region, read noise, is the random noise measured under dark conditions, generally having a slope of 0. Read noise includes a combination of pixel source follower noise, sense node reset noise, thermal dark current noise, ADC quantization noise, offset, and system noise. Read noise is overcome by photon shot noise, which on a log-log plot results in an approximate slope of  $\frac{1}{2}$ . Photon shot noise is defined as the standard deviation of the number of photon interactions per pixel, which is a



**Fig. 2** Camera under test (Canon 5D Mark III) and the integrating sphere. To ensure each pixel receives light from the whole disk-shaped illumination source (the integrating sphere output port), an f-number of 8 is required. By placing the sensor at a distance  $D$  of 81.28 cm from the output port having a diameter  $d$  of 10.16 cm, an f-number of 8 is achieved, where  $f_{\#} = d/D$



**Fig. 3** Photon Transfer Curve, illustrating the different noise regions identified as part of the overall total noise curve. Regions are defined by slopes

spatially and temporally random phenomenon. Additionally, while photons generate electron-hole pairs via the photoelectric effect, particles other than photons such as high energy electrons, protons, and ions can also produce signal carriers. Following shot noise is the region dominated by fixed pattern noise (FPN). FPN has an approximate slope of 1 indicating that signal and FPN scale together. Finally, the last region is defined as the point where the pixels enter full-well, or saturation. Full-well is defined as a rapid noise deviation from a slope of  $\frac{1}{2}$  or 1.

### 3.2 Per pixel gain factors

There are several spatially-based factors that influence the response of the imager, resulting in differences in pixel to pixel sensitivity. First, under identical light intensities, different pixels can consistently produce different DN's as a result of the manufacturing process of the imager and imperfections such as dust particles interfering with light on the sensor. The inherent differences are known as the PRNU pattern. Second, as a result of temperature variation, dark current varies between different pixels resulting in a DCNU pattern (or FPN pattern). DCNU can be corrected by subtracting a dark frame from each image frame. The dark frame is captured with the optics receiving no light intensity, and at the same temperature, exposure time, and ISO setting as the image frame. Finally, the vignette effect of the attached lens will result in variation of the light intensity for individual pixels, with the least effect generally along the optical axis and the greatest effect extending radially towards the corners of the image frame. The aperture is held constant during the measurements to minimize variation in veiling glare, which is also dependent on the lens type and manufacturing process [32].

The combined individual per-pixel gain correction accounting for spatially-based noise factors PRNU and DCNU, as well as vignette effect, is referred to as the per pixel gain  $a_j$ , defined per Granados as:

$$a_j = \frac{E[ff_j] - E[b_j]}{\frac{1}{|\Omega|} \sum_j (E[ff_j] - E[b_j])} \quad (2)$$

Per pixel gain factors are determined using the flat fielding technique [33], where bias and temporal noise are minimized via correcting for dark current by capturing both flat field  $E[ff_j]$  (uniform light field) and background frames  $E[b_j]$  (lens covered), and then averaging over 36 frames  $\Omega$ . The number of averaged frames is selected as a balance between noise reduction and data collection and processing time. Our method illuminated the sensor via the attached lens with spatially uniform light from an integrating sphere. SNR is maximized by setting the illumination such that the captured DN's are close to saturation.

### 3.3 Radiometric scene calibration

To determine the accuracy of a recombined luminance map, luminance values taken from the recombined image of two test scenes are compared with precision measurements made via a commercially calibrated luminance spot meter. Test scene 1 contains an X-Rite ColorChecker Classic as well as 6 self-printed  $3 \times 3$  grey patch charts to be used as measurement test targets throughout the scene as shown in Fig. 4. The top left, middle left, and bottom left data correspond to the

$3 \times 3$  patch charts on the left side of the scene, masked in darkness. The centre data corresponds with the middle  $3 \times 3$  chart, with the ColorChecker chart located below the centre chart. The top right and bottom right data corresponds to the  $3 \times 3$  patch charts to the right, and highest luminance, side of the scene. Absolute luminance measurements of the grey patch targets require high accuracy and sensitivity in order to properly characterize optical effects in low signal to noise regions, in this case a Konica/Minolta LS-150 luminance spot meter (accuracy 2%, sensitivity  $0.001 \frac{cd}{m^2}$ ) was used. In selecting the spot meter, the meter with the smallest available acceptance angle, 1/3 degree, was chosen in order to limit glare effects in the meter lens. A black sheet of cardboard containing a hole was experimented with to further limit meter lens glare, however positioning and robustly attaching the sheet to the meter was problematic and reduced repeatability, and therefore was not used in the measurements. The use of a cardboard mask with hole is identified as an area for future work. Luminance measurements were taken from the recombined image using the program Photosphere [29], specifically the calibration and selection functions. Photosphere luminance measurements use the radiance to average luminance conversion factor  $k$  of 179 lumens/watt, with luminance  $L$  defined in terms of RGB primaries as:

$$L = 179 * (0.265 * R + 0.760 * G + 0.065 * B). \quad (3)$$

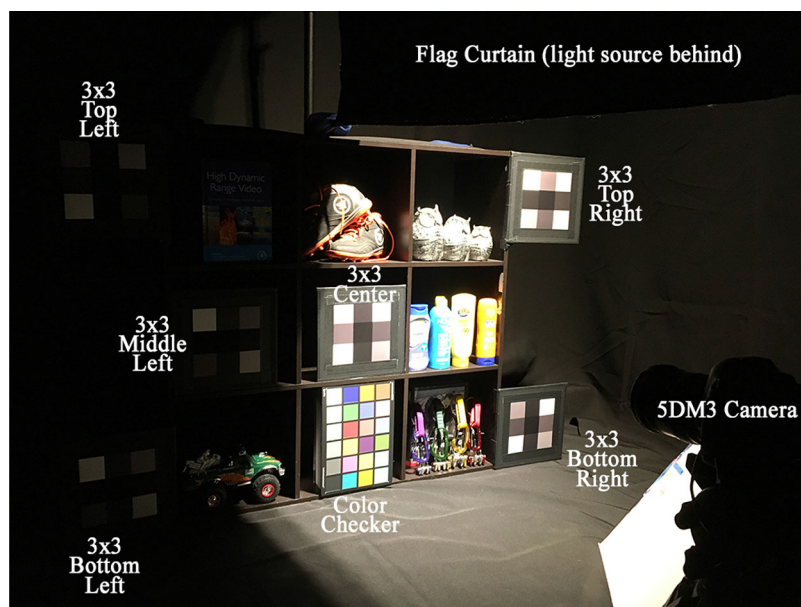
### 3.4 Glare

An ideal point light source, in focus, should illuminate only a single pixel of an image array. In practice, based on lens quality and the distance of the light source from the lens, light will be contributed to other sensor pixels as described by the glare spread function (GSF) [34]. The GSF manifests as lens flare, resulting from sharp edges internal to the lens, including the lens aperture and leaves. Additionally, at further distances lens flare becomes blurred, and the effects of stray light from lens surfaces, camera walls, and the sensor surface itself combine as veiling glare. The total achievable dynamic range of multiple exposure HDR recombination may be limited by veiling glare, potentially effecting the entire image. This can become more evident in lower luminance measurements [20].

## 4 Implementation and results

Two controlled test scenes are each lit by a quad-LED based light source (Sylvania LED9MR16/DIM/830/FL35/RP, 3000°K) powered by a regulated DC supply (Delta 1145-1071-ND). Scene 1 contains test charts and items of various specular and diffuse reflection, with the light source configured to light much of the scene, while also flagging a darker region.

**Fig. 4** HDR test scene 1 contains an X-Rite ColorChecker Classic as well as 6 self-printed 3 × 3 grey patch charts



Scene 2 is configured to specifically test the effects of glare in shadow regions, both near and far from the brighter regions of the scene. A multiple exposure set of 10 images are collected using the Canon 5DM3 plug-in MagicLantern (bracket type: exposure, centre exposure 1/8 second, 2EV step) resulting in exposure times as detailed in the Fig. 9 caption. A Sigma 24–70mm lens set at 24mm, with full open IRIS, was attached to the Canon 5DM3 and supported via tripod. The optical axis is aligned approximately with the centre of the scene.

#### 4.1 Method

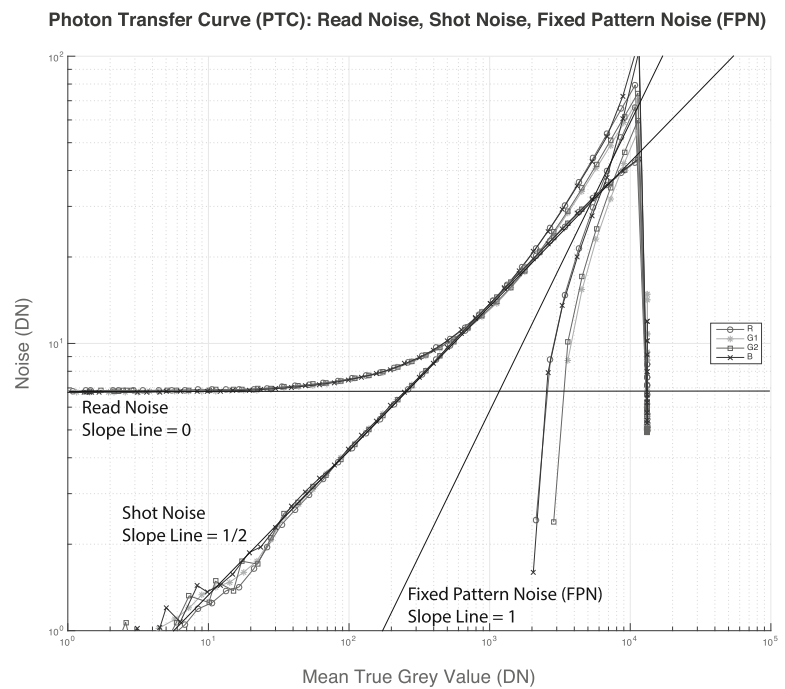
The recombination of multiple exposures includes eight key image processing steps as shown in Fig. 5, where Blocks 1–3 are single capture processes and Blocks 4–8 describe the recombination process. Figure 5, Blocks 1 and 2, illustrate the conversion of the Canon CR output .CR2 file to a .DNG file via the Adobe DNG Converter program. The Adobe converter uses ddraw as the low level converter (preference settings: Compatibility Custom, Linear unchecked, Uncompressed checked) [35]. Ddraw conversion here primarily converts the .CR2 file to a Matlab readable .DNG file. Since Canon CR data does not apply a non-linear transformation, no inverse transformation is required at this step. The active region is cropped using the Exif data, resizing from  $3950 \times 5920$  to  $3840 \times 5760$  pixels, so the resulting resolution matches that produced by commercially available software for later comparison. The individual RGGGB channels are then separated into 4 similarly sized arrays containing the individual Bayer RGGGB channels, i.e. an  $1920 \times 2880 \times 4$  array as shown in Fig. 5, block 3.



**Fig. 5** Processing Blocks illustrating the 8 key processing steps of the recombination procedure. Blocks 1–3 are single capture process and Blocks 4–8 describe the recombination process

The next step is the recombination of multiple exposures and recreation of the RGGGB Bayer pattern as shown in Fig. 5, Blocks 4 and 5. The determination of the performance parameters required for the noise-based optimal weighting function, including read noise mean and variance, charge capacity, and incident photon to ADC sensitivity (camera gain), are determined from the PTC. The FPN is determined from the flat field measurement.

**Fig. 6** Photon Transfer Curve including individual contributions of read noise, shot noise, and FPN, and linear best fit lines



**Table 1** Summary of PTC Determined Values

Performance parameter	Symbol (units)	Red Ch.	Grn Ch.1	Grn Ch.2	Blue Ch.
Incident photon to ADC sensitivity (camera gain)	$K_{ADC_{e^-/DN}}$ ( $e^-/DN$ )	7.7	5.54	5.01	4.24
Read noise mean	$\sigma_{READ_{electrons}}$ (DN)	6.84	6.78	6.84	6.75
Read noise variance	$\sigma_{R^2_{electrons}}$ (DN)	46.65	45.64	46.78	45.39
Charge capacity	$S_{FW}$ (DN)	13236	13236	13236	13236

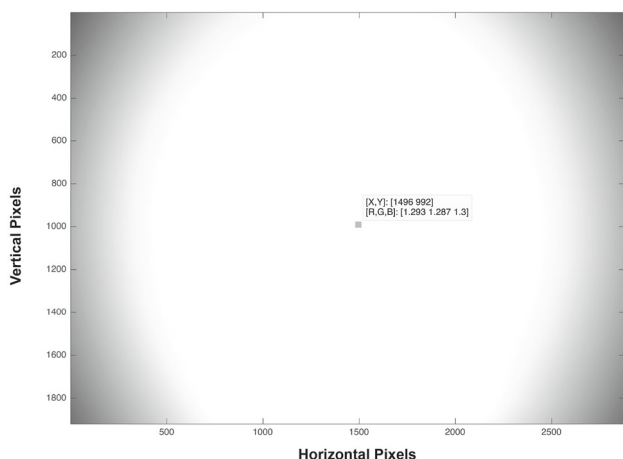
**4.1.1 PTC of the canon 5DM3**

The PTC is obtained using a grey image with exposure time of 0.005 seconds (1/200th) resulting in an approximate mid-exposure without Bayer channels nearing saturation. An iterative approach is used to determine the appropriate image sub-region meeting the uniformity requirement of 3% across the pixel sub-array (per EMVA1288), limiting error in FPN measurements. To assist in deciding the range of pixels having spatial non-uniformity of 3%, a two-dimensional finite impulse response (FIR) smoothing filter is applied to the image, smoothing the variations due to pixel noise. The PTC for the Canon 5DM3 is shown in Fig. 6, with linear slope lines for the Read Noise having a slope of 0, the Shot Noise with slope of  $\frac{1}{2}$ , and the FPN with a slope of 1. The mean Read Noise is determined from the y-axis intercept of a line with slope 0. The shot noise x-intercept of a line with slope  $\frac{1}{2}$  provides values for the incident photon to ADC sensitivities  $K_{ADC_{(e^-/DN)}}$ . Full-well (saturation) values are the same for all 4 Bayer channels. Note that in the PTC curve, the black offset (2047) is removed, resulting in a DN value for full-well of 13,236. The Exif full-well value in the captured

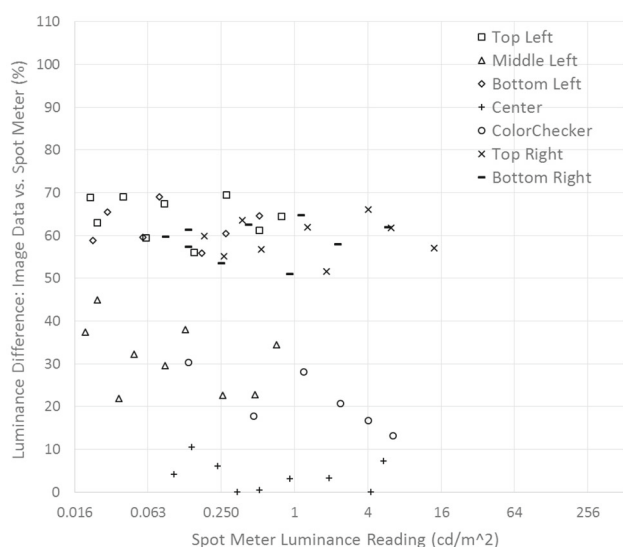
files is 15,000, whereas the PTC determined value, including black level, is  $13,236 + 2,047 = 15,283$ . The camera manufacturer specified a lower full-well value so as to introduce all channels into saturation earlier (15,000 vs. 15,283), avoiding potential colour shading or casts of the highlights. An example of colour shading is the case where the green channel saturates, with the blue and red channel just below saturation, resulting in a pink or purple highlight cast. In our processing, the PTC saturation value of 15,283 was used, so as to maintain the linearity of the original values. PTC results are summarized in Table 1.

**4.1.2 Flat fielding**

The per pixel gain correction is determined via the flat fielding process, removing PRNU, DCNU, and vignette contributions. The flat field file for the 5DM3 measurement with Sigma 24–70 mm lens is shown in Fig. 7; this is tone-mapped for visualization purposes. The per pixel contribution is observed by comparing a scene capture with and without the correction applied. The percent error graph with and without the correction applied is shown in Fig. 8. The centre  $3 \times 3$




**Fig. 7** Flat field measurement of Per Pixel Gain Factor for the Canon 5DM3 camera with Sigma 24–70 mm lens at 24 mm, full open IRIS



**Fig. 8** Luminance map percent difference with diffuse lighting, comparing with and without per pixel gain factor correction

patches are affected the least by vignette, being close to the optical centre of the lens. The ColorChecker chart vertical grey patches percent error increases with patch distance from the optical centre. The remaining charts further illustrate the increase in percentage difference at greater distances from the optical centre of the lens. We observe that the per pixel gain based error is significant, primarily a result of lens vignette, as high as 70% near the corners of the field of view. The flat fielding calibration is important not just for radiometric measurements, but also when performing test measurements of dynamic range using dynamic range test charts. Results may vary depending on where the test chart is positioned within the scene, if a flat field correction is not applied.




**Fig. 9** Scene 1 multiple exposure single images (L:R, T:B, 1/8000, 1/2000, 1/500, 1/125, 1/30, 1/8, , 2, 8, 32 s)

### 4.1.3 Recombination result

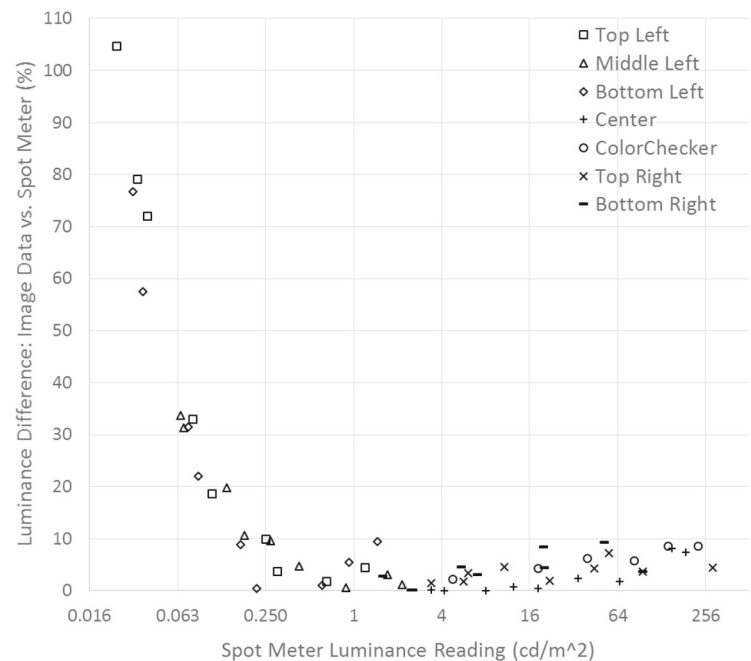
Upon combination of multiple exposures using the noise-based weighted average, the next step is to white balance the resulting data as shown in Fig. 5, Block 6. As a guide, the documentation for Coffin’s image conversion program dcrav describes different white balance behaviors concerning the highlights, known as the -H option. Modes of operation are according to the value range in which the white balance multipliers will be forced to be calculated, with a unity value indicating no correction. Multipliers less than or equal to 1 will maintain the relative proportions of unsaturated pixels, at the risk of colour artefacts in the areas that were originally saturated. In our implementation, multipliers are held to less than or equal to 1 to maintain the relative proportions. The white balance multipliers are read from the camera Exif information. In the case of the 5DM3 under test, values of 0.4756, 1.0, and 0.6611 were obtained for the colour temperature setting of 3050 K (selected to match the light source colour temperature). Demosaicing is performed upon white



**Fig. 10** Scene 1 multiple exposure recombined false colour image of Canon 5DM3 with Zeiss lens, minimum luminance set to  $0.01 \text{ cd/m}^2$  and maximum luminance set to  $1e + 05 \text{ cd/m}^2$  illustrating the linear range of the HDR image



**Fig. 11** Luminance percent difference for scene 1 grey scale patches, Canon 5DM3 with Sigma 24–70mm lens at 24 mm



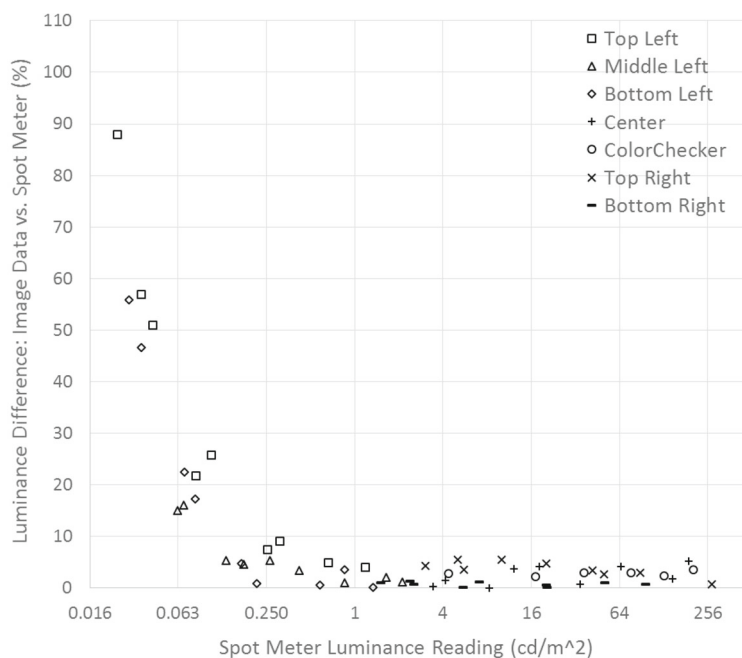
balancing, as shown in Fig. 5, Block 7. For interpolation, we implement Malvar's gradient corrected bilinear interpolated approach, while incorporating floating point values in order to maintain precision [36]. The demosaiced image is then stored in Radiance HDR format [37,38] as shown in Fig. 5, Block 8. The original image sequence single exposures are shown in Fig. 9.

#### 4.2 Effect of glare on the calibration

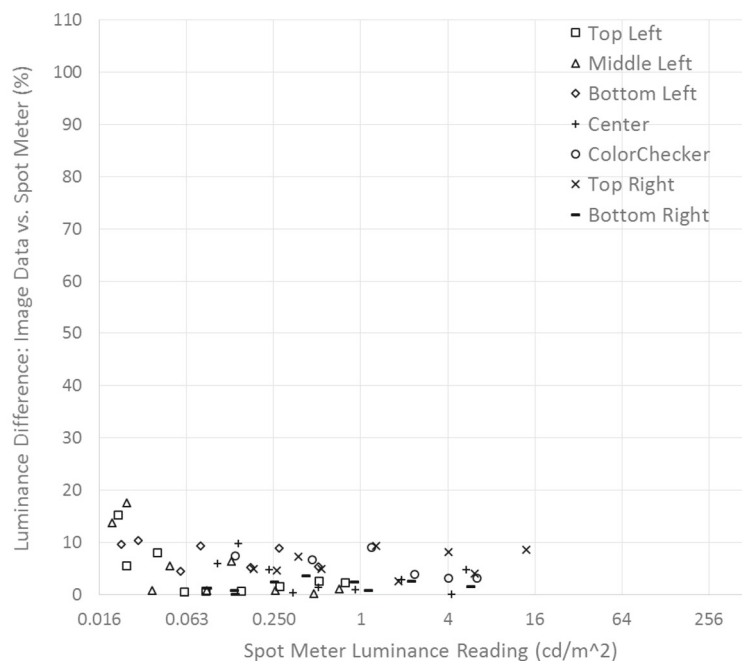
Figure 10 includes the false color representation of the resulting multiple exposure recombined image of test scene 1. Figure 11 shows percent error between the factory calibrated spot meter measurements and recombined HDR image mea-

surements for test scene 1 using a Sigma lens attached to the Canon 5DM3 camera. In Fig. 11, we observe a logarithmic percent error increase for luminance levels falling below approximately  $0.25 \text{ cd/m}^2$ , corresponding with the top left, middle left, and bottom left  $3 \times 3$  patch charts masked in darkness. The spot meter measurements were taken with a high accuracy Konica/Minolta LS-150 luminance meter, and based on this and the carefully performed HDR recombination effort, the low luminance physical accuracy error was attributed to glare. Published efforts to reduce glare, such as computational methods involving deconvolution, have not performed well as a result of low signal to noise ratio [21]. A currently recommended approach is to physically reduce glare before capturing the image, often through the use of

**Fig. 12** Luminance percent difference for scene 1 grey scale patches with Zeiss 25 mm lens



**Fig. 13** Luminance percent difference for scene 1 with diffuse lighting using Zeiss 25 mm lens



a higher quality lens, however it should be pointed out that glare free lenses due not exist. The use and comparison of a higher quality lens could provide confirmation that the observed low luminance error occurs as a consequence of glare.

**4.2.1 Influence of lens quality**

The effect of veiling glare can be illustrated via comparison of two different lens types. An additional data set is captured

using a Zeiss 25 mm lens marketed by Zeiss as “exhibiting great flare suppression through T\* anti-reflex coatings and internal light traps”. The Zeiss lens has a higher price point (currently US\$4200) than the less expensive Sigma 24–70 mm lens (\$US900). Figure 12 shows the percentage error between the spot meter measurements and the HDR image data with the Zeiss lens. The Zeiss lens results in improved glare performance as compared to the Sigma lens (Fig. 11), particularly in the low luminance range where there is an overall reduction in error. It is observed that the percent error

is now consistently under 10% from the maximum luminance patch of  $286.8 \text{ cd/m}^2$  to approximately  $0.1 \text{ cd/m}^2$ .

#### 4.2.2 Indirect (diffuse) lighting

Another method for demonstrating the non-linear low luminance response resulting from lens glare is to reduce the scene glare due to the light source. The lighting for test scene 1 is adjusted from direct illumination to reflected illumination by alternatively bouncing the light source off the white ceiling of the test environment, effectively creating a diffuse lighting environment. The diffuse lighting environment has the least effect from lens glare, providing a level of isolation to the per pixel gain factors only. Figure 13 shows the percentage error between the spot meter measurements and the HDR image data for the diffuse lighting environment using the Zeiss lens. It is evident that with the diffuse lighting, the peak luminance patch readings are reduced, measuring no more than  $16 \text{ cd/m}^2$ . More importantly, without the direct illumination as in the original scene, the percentage difference resulting from lens flare is significantly reduced, remaining under 20% through the lowest luminance levels.

#### 4.3 Spatial nature of glare

We next construct a new test scene 2 to characterize the spatial nature of glare throughout the capturing device field of view, with dark regions but near and far from the bright region. The original image sequence single exposures are shown in Fig. 14, containing a region of high luminance on the right side of the field of view, along with a dark region on the left side. A flag curtain is used to separate the sides, and white reflective cards are used to bounce similar levels of diffuse light onto the 4 dark side  $3 \times 3$  patch targets. The light side  $3 \times 3$  patches receive direct light from the lighting source. Figure 15 contains a false color version of the multiple image capture recombination of test scene 2.

Spot meter measurements are collected for each of the  $3 \times 3$  patches in the scene. Percentage difference for scene 2 between calibrated spot meter measurements, and values taken from a Canon 5DM3 image utilizing a Sigma lens at 24 mm focal length, are included in Fig. 16. The nonlinear error attributed to glare that was observed in test scene 1 for lower luminance levels (Fig. 11), is again observed in test scene 2, although in this scene the nonlinear error begins at approximately  $0.8 \text{ cd/m}^2$  as opposed to  $0.25 \text{ cd/m}^2$  of test scene 1. It is also noted that glare effects are observed on the bottom centre, bottom left, and top centre targets. Both of these points illustrate the scene and spatial dependant nature of glare.





Fig. 14 Scene 2 multiple exposure single images (L:R, T:B, 1/8000, 1/2000, 1/500, 1/125, 1/30, 1/8, 2, 8, 32s)

## 5 Discussion

### 5.1 Variability due to sensor and lens


The process of calibrating a camera sensor and lens, via the procedure described in this article, is both time intensive and expensive in terms of required laboratory equipment. We consider the reuse of camera and lens calibration data, if similar equipment is employed for image capture. Comparisons of the original Canon 5DM3 and Sigma 24–70 mm lens were made with a same model camera, and several lenses, set at the same focal length of 24 mm. The camera/lens comparison data is included in Table 2.

The lowest standard deviation was obtained using the original camera with a rented, identical make and model, Sigma 24–70 mm lens. A difference plot is shown in Fig. 17. The ring in the centre of the difference plot results from the outer radius of a port cover installed in the rear of the integrating sphere, which is used to align the two images. The borrowed 5DM3 camera (Canon 5DM3 #2), using the original Sigma



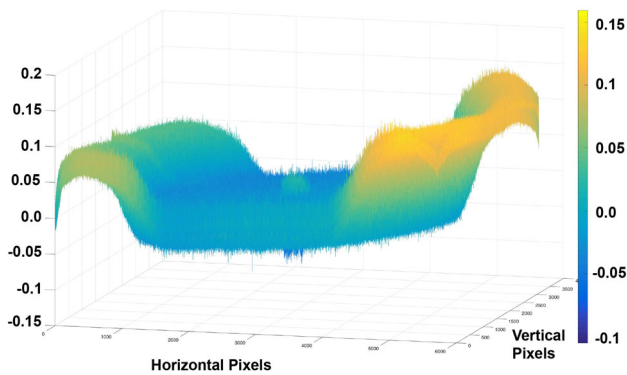
**Fig. 15** Scene 2 with dark regions both near and far from high luminance levels

**Fig. 16** Luminance percent difference for scene 2 with dark regions near and far from high luminance levels



**Table 2** Canon 5DM3 #1 with Sigma 24–70 mm #1 compared to similar equipment

Camera body	Lens (@ 24 mm)	Mean	SD
Canon 5DM3 #1	Sigma 24–70 mm #2	5.47E–04	0.0431
Canon 5DM3 #2	Sigma 24–70 mm #1	– 3.10E–03	0.1009
Canon 5DM3 #1	Canon 16–35 mm	– 5.20E–03	0.0719
Canon 5DM3 #1	Canon 24–105 mm	– 5.80E–03	0.0662



**Fig. 17** Canon 5DM3 Camera Body #1: Difference Plot with Sigma 24–70 mm Lens #1 versus Sigma 24–70 mm Lens #2

24–70 mm lens, resulted in the largest standard deviation. The last two entries are for the original camera with two additional, albeit different model, Canon lenses both set at 24 mm.

## 5.2 Variability due to image content

In scene 2 (Fig. 15), glare related error in the radiance map was determined for the bottom centre, bottom left, and top centre targets. In other words, glare may effect elements of the scene both close and further away from the high luminance areas. Of these, the bottom centre had the highest error, and when looking at the false color image, glare can be observed shaded in green seeming to “spillover” from the flag curtain towards the bottom centre target. We can deduce that any movement or changes in the light position, flag position, or other elements of the scene will result in some spatial glare variability. Predicting the effect, for variations in lighting and scene content, as well as for different optic types, is a challenging future task.

## 5.3 Comparison to commercially available software

To further assess our method, the multiple exposure CR set of 10 images (direct lighting, Zeiss lens) of scene 1 were converted to tiff images in Photoshop, and recombined into an HDR image via the “Make HDR” option in Photosphere. Figure 18 includes false colour representations of the HDR images for our method utilizing per pixel gain factor correction (Fig. 18L), and for the Photosphere method (Fig. 18R). Comparing the two images, increasing differences in luminance are observed at radially greater distances from the optical centre of the lens. This is most evident on the outer edges of the scene, coinciding with the flat field measurement illustrated in Fig. 7. Both data sets suffer from the effects of veiling glare in the low luminance regions, however the Photosphere results indicate that the lack of per pixel gain factor correction results in additional error of nearly 60% in the cor-


ner regions, when attempting to create a physically accurate radiance map.

## 5.4 Survey of methods to mitigate glare

Raskar et al. [39] statistically analyze the 4D ray-space inside a camera for a single exposure photograph. In light field terminology, 4D includes both spatial and angular variations of incident rays at each sensor location, although the method does not require reconstruction of the scene light field. While glare includes an additive low frequency bias in 2D, a significant component of glare is identified as high frequency noise in 4D ray-space that can be reduced by outlier rejection. The method requires the insertion of a high frequency mask near the camera sensor, acting as a sieve to separate spurious rays in ray-space. Results indicate glare reduction for many, but not all, types of glare, along with improvements in contrast. Rouf et al. [40] propose a computational photography approach, where highlights of a single capture standard image sensor are encoded into spatially shaped glare patterns. The low dynamic range portion of the image, and highlight information, are optically encoded using a cross screen (star) filter. The glare pattern is then used to infer the radiometric intensity distributions in the saturated image regions. Results demonstrate an observable level of reconstruction of saturated regions and removal of glare produced by the filter. The method is limited for scenes with large saturated regions, or for scenes with color gradients oriented in the same fashion as the glare pattern. Safranek [41] explored existing flare removal techniques via blind deconvolution, as well as via adaption of a high frequency occlusion mask. Deconvolution algorithms failed to provide an increase in luminance value accuracy. The high frequency occlusion mask was reported as exhibiting promising results, with limitations of light attenuation, high processing times, and the inability to apply the method to complex architectural scenes. A variant of a high frequency grid as an occlusion mask for the specific removal of veiling glare is further investigated by Cozzi et al. [42]. Specifically, the effectiveness of the mask without an estimation and subtraction when considering the non-occluded regions only. Results indicate that the use of an occlusion mask during acquisition causes a global reduction of image glare, even without estimation and subtraction. Finally, McCann and Vonikakis [43] describe a computer program for calculating the contrast image on the human retina, along with techniques for the accurate measurement of HDR scenes using superimposed photographic transparency films.

## 6 Conclusions

This article has presented a novel method of minimizing the optical effects on HDR calibration by a careful multiple



**Fig. 18** (L) Our calibrated recombination method (R) photosphere method

exposure noise-based workflow. In this new approach, the linearization of raw data from a sample camera is examined for the purpose of recombination, and creation, of the absolutely calibrated HDR scene. In processing the CR data, care is taken to account for the manufacturer defined black level and saturation limit while cropping the full sensor data to the active region. The channels of the Bayer CFA are separated and multiple exposures are combined applying an optimal HDR reconstruction based on a rigorous noise model and a noise-based weighted average. White balancing is performed while maintaining the relative proportions of the individual colour channels. The resultant demosaiced HDR image is stored as a Radiance HDR image file. The effects of per pixel gain factor and veiling glare are observed, summarized by stating that per pixel gain factor error has greater significance radially outward from the centre of the lens, and glare error has greater significance in the lower luminance regions.

Traditional methods combining multiple exposures into an HDR image often allow for absolute calibration to a known reference measurement, yet offsetting all image values based on a reference measurement does not provide a measure of the linearity of the data throughout the full captured range. In our work, the calibrated methodology identifies the lower limit of the dynamic range that falls within an acceptable absolute luminance error. The scale in the false colour image in Fig. 6 illustrates that for our captured image of scene 1, the shadow area limit is approximately  $0.1 \text{ cd/m}^2$  based on 10% acceptable error, and the peak luminance for a large enough group of pixels considered to be observable is approximately  $1e + 04 \text{ cd/m}^2$ . This results in a physically calibrated linear dynamic range of 16.6 stops.

Future work includes modeling lens glare and investigating methods to compensate for the effects as discussed in Sect. 5.4, including variations in lighting and scene content, as well as for different optic types. Also of interest is the investigation of optical effects on the use of commercially available HDR test charts and processing software to determine the dynamic range of HDR camera and lens systems. Specifically, if a per pixel gain factor correction is employed

as part of the test procedure, and if not, how the placement of the test chart within an un-calibrated optical field of view could be skewing the results. Other work will include the evaluation of luminance and chromaticity, using the colour chip section of the ColorChecker chart in test scenes. Finally, further evaluation of glare effects on the luminance spot meter lens is of interest in order to evaluate related measurement accuracy limitations.

**Acknowledgements** The authors would like to thank Dr. Jon Olansen at NASA's Johnson Space Centre for use of Morpheus sample imagery.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Arens, E.E., Youngquist, R.C., Starr, S.O.: Intensity calibrated hydrogen flame spectrum. *Int. J. Hydrog. Energy* **39**(17), 9545–9551 (2014)
2. Karr, B.A., Chalmers, A.G., Debattista, K.: High dynamic range digital imaging of spacecraft. In: Dufaux, F., Le Callet, P., Mantiuk, R., Mrak, M. (eds.) *High Dynamic Range-From Acquisition to Display and Applications*. Elsevier, Amsterdam (2016)
3. Ashok, S., Wray, B.: Predicting engine parameters using the optical spectrum of the space shuttle main engine exhaust plume. In: *10th Computing in Aerospace Conference, Computers in Aerospace Conference*. American Institute of Aeronautics and Astronautics, pp. 29–34 (1995)
4. Madzsar, G.C., National, A. United States., Space, A., Bickford, R.L., Duncan, D.B.: An overview of in-flight plume diagnostics for rocket engines, Technical report National Aeronautics and Space Administration, distributed to depository libraries in microfiche. Microfiche. [Washington, D.C. : National Aeronautics and Space Administration], 1992. 1 microfiche (1992)

5. Larson, G. W., Shakespeare, R. A.: *Rendering with Radiance The Art and Science of Lighting Visualization*. Morgan Kaufmann Publishers, Burlington (1998)
6. Reinhard, C.F., Herkel, S.: The simulation of annual daylight illuminance distributions—a state-of-the-art comparison of six radiance-based methods. *Energy Build.* **32**(2), 167–187 (2000)
7. Mantiuk, R., Myszkowski, K., Seidel, H.-P.: A perceptual framework for contrast processing of high dynamic range images. *ACM Trans. Appl. Percept.* **3**(3), 286–308 (2006)
8. Akyüz, A.O., Reinhard, E.: Noise reduction in high dynamic range imaging. *J. Vis. Commun. Image Represent.* **18**(5), 366–376 (2007)
9. Olansen, J.B.: Project morpheus: lander technology development. In: *AIAA SPACE 2014 Conference and Exposition*, p. 4314 (2014)
10. Karr, B., Debattista, K., Chalmers, A.: Chapter 4—calibrated measurement of imager dynamic range. In: Chalmers, A., Campisi, P., Shirley, P., Olaizola, I.G. (eds.) *High Dynamic Range Video*, pp. 87–108. Academic Press, Cambridge (2017)
11. DiCarlo, J.M., Wandell, B.A.: Rendering high dynamic range images. In: *Electronic Imaging, International Society for Optics and Photonics*, pp. 392–401 (2000)
12. Mann, S., Picard, R.W.: On being “undigital” with digital cameras: extending dynamic range by combining differently exposed pictures. In: *proceedings of IS&T 48th Annual Conference Society for Imaging Science and Technology Annual Conference* (1995)
13. Mitsunaga, T., Nayar, S.K.: Radiometric self calibration. In: *IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 1999*, vol. 1, IEEE (1999)
14. Tsin, Y., Ramesh, V., Kanade, T.: Statistical calibration of ccd imaging process. In: *Eighth IEEE International Conference on Computer Vision, 2001. ICCV 2001. Proceedings*, vol. 1, pp. 480–487 (2001)
15. Granados, M., Ajdin, B., Wand, M., Theobalt, C., Seidel, H.-P., Lensch, H. P.: Optimal hdr reconstruction with linear digital cameras. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2010 IEEE*, pp. 215–222 (2010)
16. Griffiths, D.J., Wicks, A.: Simplified weighting function for high dynamic range video frame formation. *Appl. Opt.* **55**(31), C9–C17 (2016)
17. Hasinoff, S.W., Durand, F., Freeman, W.T.: Noise-optimal capture for high dynamic range photography. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2010*, pp. 553–560 (2010)
18. Hasinoff, S.W., Sharlet, D., Geiss, R., Adams, A., Barron, J.T., Kainz, F., Chen, J., Levoy, M.: Burst photography for high dynamic range and low-light imaging on mobile cameras. *ACM Trans. Graph.* **35**(6), 192 (2016)
19. Janesick, J.: *Photon Transfer DN to Lambda*. SPIE Press, Bellingham (2007)
20. McCann, J.J., Rizzi, A.: Camera and visual veiling glare in hdr images. *J. Soc. Inf. Display* **15**(9), 721–730 (2007)
21. Talvala, E.-V., Adams, A., Horowitz, M., Levoy, M.: Veiling glare in high dynamic range imaging. *ACM Trans. Graph.* **26**, 37 (2007)
22. ISO, Noise measurements. electronic still-picture imaging. photography., Standard BS ISO 15739:2013, BSI Standards Limited (2013)
23. Debevec, P.E., Malik, J.: Recovering high dynamic range radiance maps from photographs. In: *ACM SIGGRAPH 2008 classes*, pp. 1–10 (2008)
24. Inanici, M.N., Galvin, J.: Evaluation of high dynamic range photography as a luminance data acquisition system. *Light. Res. Technol.* **38**(2), 123–136 (2004)
25. Coutelier, B., Oumortier, O.: Luminance calibration of the nikon coolpix 990 digital camera. *Publications-Commission Internationale De L Eclairage Cie* **152**, D3–56 (2003)
26. Robertson, M. A., Borman, S., Stevenson, R. L.: Dynamic range improvement through multiple exposures. In: *Proceedings 1999 International Conference on Image Processing (Cat. 99CH36348)*, Vol. 3, pp. 159–163 (1999)
27. Reinhard, E., Heidrich, W., Debevec, P., Pattanaik, S., Ward, G., Myszkowski, K.: *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*. Elsevier, Amsterdam (2010)
28. Kirk, K., Andersen, H. J.: Noise characterization of weighting schemes for combination of multiple exposures. In: *BMVC*, vol. 3, pp. 1129–1138 (2006)
29. Ward, G.: *Photosphere software archive* (2010)
30. Popadić, I., Todorović, B. M.: Image fusion based on the multi-exposed images. In: *2016 24th Telecommunications Forum (TELFOR)*, pp. 1–4 (2016)
31. EMVA.: *Standard for characterization of image sensors and cameras*, Technical Report EMVA Standard 1288, European Machine Vision Association Nov 29 (2010)
32. Tomic, I., Karlovic, I., Juric, I.: Practical assessment of veiling glare in camera lens system. *J. Graph. Eng. Des.* **5**, 23–28 (2014)
33. Janesick, J.: *Scientific Charge-Coupled Devices*. SPIE Publications, Bellingham (2001)
34. ISO, Optics and optical instruments, veiling glare of image forming systems, definitions and methods of measurement, Standard ISO9358:1994, International Organization for Standardization (1994)
35. Coffin, D.: *Decoding raw digital photos in linux* (2017)
36. Malvar, H. S., He, L.-w., Cutler, R.: High-quality linear interpolation for demosaicing of bayer-patterned color images. In: *Acoustics, Speech, and Signal Processing, 2004. Proceedings.(ICASSP'04)*. IEEE International Conference on, Vol. 3, IEEE, pp. iii–485 (2004)
37. Ward, G.: Real pixels. In: Arvo, J. (ed.) *Graphics Gems II*, vol. 2, pp. 80–83. Academic Press, Cambridge (1991)
38. Larson, G.W.: *High dynamic range image encodings*. Technical Report (2005)
39. Raskar, R., Agrawal, A., Wilson, C.A., Veeraraghavan, A.: Glare aware photography: 4d ray sampling for reducing glare effects of camera lenses. *ACM Trans. Graph.* **27**(3), 1–10 (2008)
40. Rouf, M., Mantiuk, R., Heidrich, W., Trentacoste, M., Lau, C.: Glare encoding of high dynamic range images. In: *CVPR 2011*, pp. 289–296 (2011)
41. Safranek, S. F.: A comparison of techniques used for the removal of lens flare found in high dynamic range luminance measurements
42. Federico, C., Carmine, E., Giovanni, G., Filippo, R., Matteo, L., Alessandro, R.: Use of an occlusion mask for veiling glare removal in hdr images. *J. Imaging* **48**, 100 (2018)
43. John, J.M., Vassilios, V.: Calculating retinal contrast from scene content: a program. *Front. Psychol.* **8**, 2079 (2018)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Brian Karr** received a B.S. degree in electrical engineering from Rochester Institute of Technology, Rochester, NY, 1993, a M.S. degree in electrical engineering from University of Central Florida, Orlando, FL, 1998, and a Ph.D. degree in electrical engineering from the University of Warwick, Coventry, UK, 2019. His research interests include image collection, processing, and analysis techniques for high dynamic range imaging.



**Kurt Debattista** is an Associate Professor within the Visualization Group, WMG, University of Warwick. He has a Ph.D. from the University of Bristol. His research interests include high-fidelity rendering, perceptual imaging, high-dynamic-range imaging, and high-performance computing.



**Alan Chalmers** is Professor, head of the Visualization Group, University of Warwick, UK and a former Royal Society Industrial Fellow. He has an M.Sc. with distinction from Rhodes University, 1985 and a Ph.D. from University of Bristol, 1991. He is Honorary President of Afrigraph and a former Vice President of ACM SIGGRAPH. Chalmers has published over 250 papers in journals and international conferences on HDR, high-fidelity virtual environments, multi-sensory perception, parallel processing and virtual archaeology and successfully supervised 47 Ph.D. students. In addition, Chalmers is a UK representative on IST/37 considering standards within MPEG.